Markov ⇔ Kalman Filter Localization

- Markov localization
  - *localization starting from any unknown position*
  - *recovers from ambiguous situation.*
  - *However, to update the probability of all positions within the whole state space at any time requires a discrete representation of the space (grid). The required memory and calculation power can thus become very important if a fine grid is used.*
- Kalman filter localization
  - *tracks the robot and is inherently very precise and efficient.*
  - *However, if the uncertainty of the robot becomes to large (e.g. collision with an object) the Kalman filter will fail and the position is definitively lost.*

Markov Localization (1)

- Markov localization uses an explicit, discrete representation for the probability of all position in the state space.
- This is usually done by representing the environment by a grid or a topological graph with a finite number of possible states (positions).
- During each update, the probability for each state (element) of the entire space is updated.

Markov Localization (2): Applying probabilty theory to robot localization

- P(A): Probability that A is true.
  - *e.g.* $p(r_t = l)$: *probability that the robot r is at position l at time t (prior)*

- We wish to compute the probability of each individual robot position given actions and sensor measures.
- P(A|B): Conditional probability of A given that we know B.
  - *e.g.* $p(r_t = l|\ i_t)$: *probability that the robot is at position l given the sensors input* i_t.
- Product rule: $$p(A \wedge B) = p(A|B)p(B)$$ $$p(A \wedge B) = p(B|A)p(A)$$
- Bayes rule: $$p(A|B) = \frac{p(B|A)p(A)}{p(B)}$$

Bayes rule example

- Suppose a robot obtains measurement *z*
- What is *P(open|z)?*

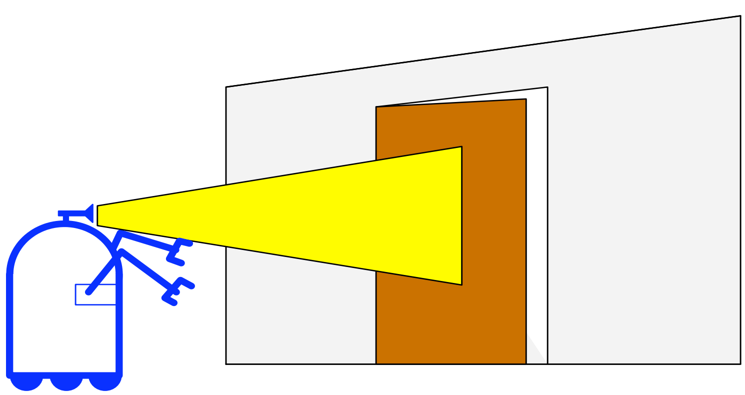

4

Bayes rule Causal vs. Diagnostic Reasoning

- *P(open|z)* is diagnostic.
- *P(z|open)* is causal.
- Often causal knowledge is easier to obtain.
- Bayes rule allows us to use causal knowledge:

count frequencies!

$$P(open \mid z) = \frac{P(z \mid open)P(open)}{P(z)}$$

5

Bayes Formula aganin

$$P(x, y) = P(x \mid y)P(y) = P(y \mid x)P(x)$$

$$\Rightarrow$$

$$P(x \mid y) = \frac{P(y \mid x)\ P(x)}{P(y)} = \frac{\text{likelihood} \cdot \text{prior}}{\text{evidence}}$$

Example

- Consider that we have a random discrete random variable X, with two outcomes characterizing whether the door is open or not. Our initial belief (or prior probability) is

  $bel(X = open) = 0.5$ and $bel(X = closed) = 0.5$

- Now suppose that we have some noisy sensors trying make some measurements of the door. The characteristics of the sensors obtained in the training/learning stage are following – the sensor can have two outcomes, each with the following conditional probability

$$P(z = sense\_open \mid X = open) = 0.6$$
$$P(z = sense\_closed \mid X = open) = 0.4$$
$$P(z = sense\_open \mid X = closed) = 0.3$$
$$P(z = sense\_closed \mid X= closed) = 0.7$$

- This suggests that detecting closed door is relatively reliable (only 0.3 errors), but detecting open door is less reliable

Example

- Suppose you get a measurement $z = sense\_open$ now
- You want to compute probability that the door is open
- To shorten the notation :
- $P(z|open) = 0.6$ $\quad$ $P(z|\neg open) = 0.3$
- $P(open) = P(\neg open) = 0.5$

$$P(open|z) = \frac{P(z \mid open)P(open)}{P(z \mid open)p(open) + P(z \mid \neg open)p(\neg open)}$$

$$P(open|z) = \frac{0.6 \cdot 0.5}{0.6 \cdot 0.5 + 0.3 \cdot 0.5} = \frac{2}{3} = 0.67$$

- ***z raises the probability that the door is open.***

8

Combining Evidence

- Suppose our robot obtains another observation z_2, from other sensing modality which has slightly different conditional probabilities
- $P(z_2|open) = 0.5$ $P(z_2|\neg open) = 0.6$
- Which for example is more detail is

$P(z_2 = sense\_open \,|X = open) = 0.5$, $P(z_2 = sense\_open \,|X = \neg open) = 0.6$

$P(z_2 = sense\_closed \,|X = open) = 0.5$, $P(z_2 = sense\_closed \,|X = \neg open) = 0.4$

- *Given these conditional prob. and knowing the prior, you can compute the joint distribution (of z_2 and X and convince your self that all entries sum up to 1*
- How can we integrate this new information?
- More generally, how can we estimate $P(x|\, z_1 ... z_n)$?

9

Recursive Bayesian Updating

$$P(x \mid z_1, \ldots, z_n) = \frac{P(z_n \mid x, z_1, \ldots, z_{n-1})\, P(x \mid z_1, \ldots, z_{n-1})}{P(z_n \mid z_1, \ldots, z_{n-1})}$$

Markov assumption:* z_n *is independent of* $z_1, \ldots, z_{n-1}$ *if we know x.

$$P(x \mid z_1, \ldots, z_n) = \frac{P(z_n \mid x)\, P(x \mid z_1, \ldots, z_{n-1})}{P(z_n \mid z_1, \ldots, z_{n-1})}$$

$$= \eta\, P(z_n \mid x)\, P(x \mid z_1, \ldots, z_{n-1})$$

$$= \eta_{1 \ldots n} \prod_{i=1 \ldots n} P(z_i \mid x)\, P(x)$$

10

Example: Second Measurement

- *$P(z_2|open) = 0.5$* *$P(z_2|\neg open) = 0.6$*
- *$P(open|z_1)=2/3$*

$$P(open \mid z_2, z_1) = \frac{P(z_2 \mid open)\, P(open \mid z_1)}{P(z_2 \mid open)\, P(open \mid z_1) + P(z_2 \mid \neg open)\, P(\neg open \mid z_1)}$$

$$= \frac{\frac{1}{2} \cdot \frac{2}{3}}{\frac{1}{2} \cdot \frac{2}{3} + \frac{3}{5} \cdot \frac{1}{3}} = \frac{5}{8} = 0.625$$

- z_2 lowers the probability that the door is open.

11

Actions

- Often the world is **dynamic** since
  - ***actions carried out by the robot,***
  - ***actions carried out by other agents,***
  - *or just the* ***time*** *passing by*

  change the world.

- How can we **incorporate** such **actions**?
- How to model them explicitely ?

Typical Actions

- The robot **turns its wheels** to move
- The robot **uses its manipulator** to grasp an object
- Plants grow over **time**…

- Actions are **never carried out with absolute certainty**.
- In contrast to measurements, **actions generally increase the uncertainty**.

Modeling Actions

- To incorporate the outcome of an action *u* into the current "belief", we use the conditional pdf

$$P(x|u,x')$$

- This term specifies the pdf that **executing *u* changes the state from *x' to x***.

State Transitions

P(x|u,x') for *u* = "close door":

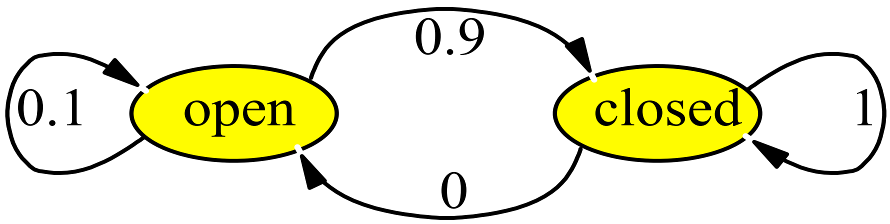

If the door is open, the action "close door" succeeds in 90% of all cases.

15

Integrating the Outcome of Actions

Continuous case:

$$P(x \mid u) = \int P(x \mid u, x')P(x')dx'$$

Discrete case:

$$P(x \mid u) = \sum P(x \mid u, x')P(x')$$

Given and action u, what will happen to the posterior Belief about state

16

Example: The Resulting Belief

$$P(closed \mid u) = \sum P(closed \mid u, x')P(x')$$
$$= P(closed \mid u, open)P(open)$$
$$+ P(closed \mid u, closed)P(closed)$$
$$= \frac{9}{10} * \frac{5}{8} + \frac{1}{1} * \frac{3}{8} = \frac{15}{16}$$
$$P(open \mid u) = \sum P(open \mid u, x')P(x')$$
$$= P(open \mid u, open)P(open)$$
$$+ P(open \mid u, closed)P(closed)$$
$$= \frac{1}{10} * \frac{5}{8} + \frac{0}{1} * \frac{3}{8} = \frac{1}{16}$$
$$= 1 - P(closed \mid u)$$

17

General Bayes Filters: Framework

- **Putting actions and observations together**
- **Given:**
  - *Stream of observations z and action data u:*

$$d_t = \{u_1, z_1 \ldots, u_t, z_t\}$$

  - *Sensor model P(z|x) (actual models discussed later)*
  - *Action model P(x|u,x') (actual models discussed later)*
  - *Prior probability of the system state P(x).*
- **Wanted:**
  - *Estimate of the state X of a dynamical system.*
  - *The posterior of the state is also called **Belief**:*

$$Bel(x_t) = P(x_t \mid u_1, z_1 \ldots, u_t, z_t)$$

18

Markov Assumption

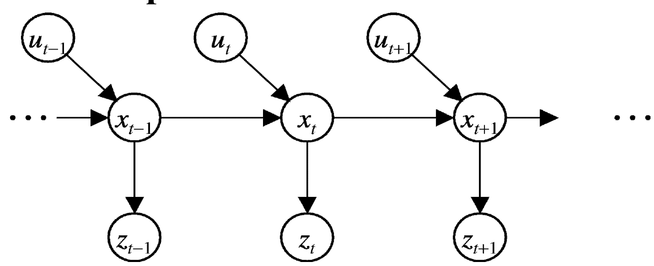

$$p(z_t \mid x_{0:t}, z_{1:t}, u_{1:t}) = p(z_t \mid x_t)$$
$$p(x_t \mid x_{1:t-1}, z_{1:t}, u_{1:t}) = p(x_t \mid x_{t-1}, u_t)$$

Underlying Assumptions

- Static world
- Independent noise
- Perfect model, no approximation errors

19

Bayes Filters

***z* = observation**
***u* = action**
***x* = state**

$$Bel(x_t) = P(x_t \mid u_1, z_1 \ldots, u_t, z_t)$$

Bayes $$= \eta \; P(z_t \mid x_t, u_1, z_1, \ldots, u_t) \; P(x_t \mid u_1, z_1, \ldots, u_t)$$

Markov $$= \eta \; P(z_t \mid x_t) \; P(x_t \mid u_1, z_1, \ldots, u_t)$$

Total prob. $$= \eta \; P(z_t \mid x_t) \int P(x_t \mid u_1, z_1, \ldots, u_t, x_{t-1}) \; P(x_{t-1} \mid u_1, z_1, \ldots, u_t) \; dx_{t-1}$$

Markov $$= \eta \; P(z_t \mid x_t) \int P(x_t \mid u_t, x_{t-1}) \; P(x_{t-1} \mid u_1, z_1, \ldots, u_t) \; dx_{t-1}$$

Markov $$= \eta \, P(z_t \mid x_t) \int P(x_t \mid u_t, x_{t-1}) \, P(x_{t-1} \mid u_1, z_1, \ldots, z_{t-1}) \; dx_{t-1}$$

$$= \eta \; P(z_t \mid x_t) \int P(x_t \mid u_t, x_{t-1}) \; Bel(x_{t-1}) \; dx_{t-1}$$

20

$$Bel(x_t) = \eta \ P(z_t \mid x_t) \int P(x_t \mid u_t, x_{t-1}) \ Bel(x_{t-1}) \ dx_{t-1}$$

Algorithm **Bayes_filter**(*Bel(x),d*):

η=0

If *d* is a perceptual data item *z* then

For all *x* do

$$Bel'(x) = P(z \mid x) Bel(x)$$

$$\eta = \eta + Bel'(x)$$

For all *x* do

$$Bel'(x) = \eta^{-1} Bel'(x)$$

Else if *d* is an action data item *u* then

For all *x* do

$$Bel'(x) = \int P(x \mid u, x') \ Bel(x') \ dx'$$

Return *Bel'(x)*

Bayes Filters are Familiar!

$$Bel(x_t) = \eta \ P(z_t \mid x_t) \int P(x_t \mid u_t, x_{t-1}) \ Bel(x_{t-1}) \ dx_{t-1}$$

- Kalman filters
- Particle filters
- Hidden Markov models
- Dynamic Bayesian networks
- Partially Observable Markov Decision Processes (POMDPs)

Summary

- Bayes rule allows us to compute probabilities that are hard to assess otherwise.
- Under the Markov assumption, recursive Bayesian updating can be used to efficiently combine evidence.
- Bayes filters are a probabilistic tool for estimating the state of dynamic systems.

23